

Measure Concentration of Markov Tree Processes

Leonid Kontorovich
 School of Computer Science
 Carnegie Mellon University
 Pittsburgh, PA 15213
 USA
 lkontor@cs.cmu.edu

February 2, 2008

Abstract

We prove an apparently novel concentration of measure result for Markov tree processes. The bound we derive reduces to the known bounds for Markov processes when the tree is a chain, thus strictly generalizing the known Markov process concentration results. We employ several techniques of potential independent interest, especially for obtaining similar results for more general directed acyclic graphical models.

1 Introduction

An emerging paradigm for proving concentration results for nonproduct measures is to quantify the dependence between the variables and state the bounds in terms of that dependence (see [3] for an overview). A process (measure) particularly amenable to this approach is the Markov process. Using different techniques, Marton (coupling method [6], 1996), Samson (log-Sobolev inequality [8], 2000) and Kontorovich and Ramanan (martingale differences [3], 2006) have obtained qualitatively similar concentration of measure results for Markov processes. One natural generalization of the Markov process is the *hidden Markov process*; we proved a concentration result for this class in [2]. A different way to generalize the Markov process is via the *Markov tree process*, which we address in the present paper.

If (\mathcal{S}^n, d) is a metric space and $(X_i)_{1 \leq i \leq n}$, $X_i \in \mathcal{S}$ is a random process, a measure concentration result (for the purposes of this paper) is an inequality stating that for any 1-Lipschitz (with respect to d) function $f : \mathcal{S}^n \rightarrow \mathbb{R}$, we have

$$\mathbb{P}\{|f(X) - \mathbb{E}f(X)| > t\} \leq 2 \exp(-Kt^2), \quad (1)$$

where K may depend on n but not on f .¹

The quantity $\bar{\eta}_{ij}$, defined below, has proved useful for obtaining concentration results. For $1 \leq i < j \leq n$, $y \in \mathcal{S}^{i-1}$ and $w \in \mathcal{S}$, let

$$\mathcal{L}(X_j^n | X_1^{i-1} = y, X_i = w)$$

be the law of X_j^n conditioned on $X_1^{i-1} = y$ and $X_i = w$. Define

$$\eta_{ij}(y, w, w') = \|\mathcal{L}(X_j^n | X_1^{i-1} = y, X_i = w) - \mathcal{L}(X_j^n | X_1^{i-1} = y, X_i = w')\|_{\text{TV}} \quad (2)$$

¹See [5] for a much more general notion of concentration.

and

$$\bar{\eta}_{ij} = \sup_{y \in \mathcal{S}^{i-1}} \sup_{w, w' \in \mathcal{S}} \eta_{ij}(y, w, w')$$

where $\|\cdot\|_{\text{TV}}$ is the total variation norm (see §2.1 to clarify notation).

Let Γ and Δ be upper-triangular $n \times n$ matrices, with $\Gamma_{ii} = \Delta_{ii} = 1$ and

$$\Gamma_{ij} = \sqrt{\bar{\eta}_{ij}}, \quad \Delta_{ij} = \bar{\eta}_{ij}$$

for $1 \leq i < j \leq n$.

For the case where $\mathcal{S} = [0, 1]$ and d is the Euclidean metric on \mathbb{R}^n , Samson [8] showed that if $f : [0, 1]^n \rightarrow \mathbb{R}$ is convex and Lipschitz with $\|f\|_{\text{Lip}} \leq 1$, then

$$\mathbb{P}\{|f(X) - \mathbb{E}f(X)| > t\} \leq 2 \exp\left(-\frac{t^2}{2\|\Gamma\|_2^2}\right) \quad (3)$$

where $\|\Gamma\|_2$ is the ℓ_2 operator norm of the matrix Γ ; Marton [7] has a comparable result.

For the case where \mathcal{S} is countable and d is the (normalized) Hamming metric on \mathcal{S}^n ,

$$d(x, y) = \frac{1}{n} \sum_{i=1}^n \mathbb{1}_{\{x_i \neq y_i\}},$$

Kontorovich and Ramanan [3] showed that if $f : \mathcal{S}^n \rightarrow \mathbb{R}$ is Lipschitz with $\|f\|_{\text{Lip}} \leq 1$, then

$$\mathbb{P}\{|f(X) - \mathbb{E}f(X)| > t\} \leq 2 \exp\left(-\frac{nt^2}{2\|\Delta\|_\infty^2}\right) \quad (4)$$

where $\|\Delta\|_\infty$ is the ℓ_∞ operator norm of the matrix Δ , also given by

$$\|\Delta\|_\infty = \max_{1 \leq i < n} (1 + \bar{\eta}_{i,i+1} + \dots + \bar{\eta}_{i,n}). \quad (5)$$

This leads to a strengthening of the Markov measure concentration result in Marton [6].

The sharpest currently known Markov measure concentration results were obtained in [3] and [8], in terms of the contraction coefficients $(\theta_i)_{1 \leq i < n}$ of the Markov process:

$$\bar{\eta}_{ij} \leq \theta_i \theta_{i+1} \dots \theta_{j-1}. \quad (6)$$

In this paper, we prove a bound on $\bar{\eta}_{ij}$ in terms of the contraction coefficients of the Markov tree process (Theorem 2.1). This bound is cumbersome to state without preliminary definitions, but it reduces to (6) in the case where the Markov tree is a chain.

2 Bounding $\bar{\eta}_{ij}$ for Markov tree processes

2.1 Notational preliminaries

Random variables are capitalized (X), specified state sequences are written in lowercase (x), the shorthand $X_i^j \equiv X_i \dots X_j$ is used for all sequences, and the concatenation of the sequences x and y is denoted by xy , as in $x_i^j x_{j+1}^k = x_i^k$. Another way to index collections of variables is by subset: if $I = \{i_1, i_2, \dots, i_m\}$ then we write $x_I \equiv x[I] \doteq \{x_{i_1}, x_{i_2}, \dots, x_{i_m}\}$; we will write x_I and $x[I]$ interchangeably, as dictated by convenience. To avoid cumbersome subscripts, we will also occasionally use the bracket notation for vector components. Thus, $\mathbf{u} \in \mathbb{R}^{\mathcal{S}^I}$, then

$$\mathbf{u}_{x_I} \equiv \mathbf{u}_{x[I]} \equiv \mathbf{u}[x[I]] \doteq \mathbf{u}_{(x_{i_1}, x_{i_2}, \dots, x_{i_m})} \in \mathbb{R}$$

for each $x[I] \in \mathcal{S}^I$. A similar bracket notation will apply for matrices.

We will use $|\cdot|$ to denote set cardinalities. Sums will range over the entire space of the summation variable; thus $\sum_{x_i^j} f(x_i^j)$ stands for $\sum_{x_i^j \in \mathcal{S}^{j-i+1}} f(x_i^j)$, and $\sum_{x[I]} f(x[I])$ is shorthand for $\sum_{x[I] \in \mathcal{S}^I} f(x[I])$.

The probability operator $\mathbb{P}\{\cdot\}$ is defined with respect the measure space specified in context.

We will write $[n]$ for the set $\{1, \dots, n\}$. Anytime $\|\cdot\|$ appears without a subscript, it will always denote the **total variation** norm $\|\cdot\|_{\text{TV}}$, which we define here, for any signed measure τ on a countable set \mathcal{X} , by

$$\|\tau\|_{\text{TV}} \doteq \frac{1}{2} \sum_{x \in \mathcal{X}} |\tau(x)|. \quad (7)$$

If $G = (V, E)$ is a graph, we will frequently abuse notation and write $u \in G$ instead of $u \in V$, blurring the distinction between a graph and its vertex set. This notation will carry over to set-theoretic operations ($G = G_1 \cap G_2$) and indexing of variables (e.g., X_G).

Unless we will need to refer explicitly to a σ -algebra, we will suppress it in the probability space notation, using less rigorous formulations, such as “Let μ be a measure on \mathcal{S}^n ”. Furthermore, to avoid the technical but inessential complications associated with infinite sets, we will take \mathcal{S} to be finite in this paper, noting only that the bounds carry over unchanged to the countable case (as done in [3] and [2]). To extend the results to the continuous case, some mild measure-theoretic assumptions are needed (see [7]).

2.2 Definition of Markov tree process

2.2.1 Graph-theoretic preliminaries

Consider a directed acyclic graph $G = (V, E)$, and define a partial order \prec_G on G by the transitive closure of the relation

$$u \prec_G v \quad \text{if} \quad (u, v) \in E.$$

We define the **parents** and **children** of $v \in V$ in the natural way:

$$\text{parents}(v) = \{u \in V : (u, v) \in E\}$$

and

$$\text{children}(v) = \{w \in V : (v, w) \in E\}.$$

If G is connected and each $v \in V$ has at most one parent, G is called a **(directed) tree**. In a tree, whenever $u \prec_G v$ there is a unique directed path from u to v . A tree T always has a unique minimal (w.r.t. \prec_T) element $r_0 \in V$, called its **root**. Thus, for every $v \in V$ there is a unique directed path $r_0 \prec_T r_1 \prec_T \dots \prec_T r_d = v$; define the **depth** of v , $\text{dep}_T(v) = d$, to be the length (i.e., number of edges) of this path. Note that $\text{dep}_T(r_0) = 0$. We define the depth of the tree by $\text{dep}(T) = \sup_{v \in T} \text{dep}_T(v)$.

For $d = 0, 1, \dots$ define the d^{th} **level** of the tree T by

$$\text{lev}_T(d) = \{v \in V : \text{dep}_T(v) = d\};$$

note that the levels induce a disjoint partition on V :

$$V = \bigcup_{d=1}^{\text{dep}(T)} \text{lev}_T(d).$$

We define the **width** of a tree as the greatest number of nodes in any level:

$$\text{wid}(T) = \sup_{1 \leq d \leq \text{dep}(T)} |\text{lev}_T(d)|. \quad (8)$$

We will consistently take $|V| = n$ for finite V . An ordering $J : V \rightarrow \mathbb{N}$ of the nodes is said to be **breadth-first** if

$$\text{dep}_T(u) < \text{dep}_T(v) \implies J(u) < J(v). \quad (9)$$

Since every directed tree $T = (V, E)$ has some breadth-first ordering,² we shall henceforth blur the distinction between $v \in V$ and $J(v)$, simply taking $V = [n]$ (or $V = \mathbb{N}$) and assuming that $\text{dep}_T(u) < \text{dep}_T(v) \Rightarrow u < v$ holds. This will allow us to write \mathcal{S}^V simply as \mathcal{S}^n for any set \mathcal{S} .

Note that we have two orders on V : the partial order \prec_T , induced by the tree topology, and the total order $<$, given by the breadth-first enumeration. Observe that $i \prec_T j$ implies $i < j$ but not vice versa.

If $T = (V, E)$ is a tree and $u \in V$, we define the **subtree** induced by u , $T_u = (V_u, E_u)$ by $V_u = \{v \in V : u \preceq_T v\}$, $E_u = \{(v, w) \in E : v, w \in V_u\}$.

2.2.2 Markov tree measure

If \mathcal{S} is a finite set, a **Markov tree measure** μ is defined on \mathcal{S}^n by a tree $T = (V, E)$ and transition kernels p_0 , $\{p_{ij}(\cdot | \cdot)\}_{(i,j) \in E}$. Continuing our convention in §2.2.1, we have a breadth-first order $<$ and the total order \prec_T on V , and take $V = \{1, \dots, n\}$. Together, the topology of T and the transition kernels determine the measure μ on \mathcal{S}^n :

$$\mu(x) = p_0(x_1) \prod_{(i,j) \in E} p_{ij}(x_j | x_i). \quad (10)$$

A measure on \mathcal{S}^n satisfying (10) for some T and $\{p_{ij}\}$ is said to be **compatible** with tree T ; a measure is a Markov tree measure if it is compatible with some tree.

Suppose \mathcal{S} is a finite set and $(X_i)_{i \in \mathbb{N}}$, $X_i \in \mathcal{S}$ is a random process defined on $(\mathcal{S}^\mathbb{N}, \mathbb{P})$. If for each $n > 0$ there is a tree $T^{(n)} = ([n], E^{(n)})$ and a Markov tree measure μ_n compatible with $T^{(n)}$ such that for all $x \in \mathcal{S}^n$ we have

$$\mathbb{P}\{X_1^n = x\} = \mu_n(x)$$

then we call X a **Markov tree process**. The trees $\{T^{(n)}\}$ are easily seen to be consistent in the sense that $T^{(n)}$ is an induced subgraph of $T^{(n+1)}$. So corresponding to any Markov tree process is the unique infinite tree $T = (\mathbb{N}, E)$. The uniqueness of T is easy to see, since for $v > 1$, the parent of v is the smallest $u \in \mathbb{N}$ such that

$$\mathbb{P}\{X_v = x_v | X_1^u = x_1^u\} = \mathbb{P}\{X_v = x_v | X_u = x_u\};$$

thus \mathbb{P} determines the topology of T .

It is straightforward to verify that a Markov tree process $\{X_v\}_{v \in T}$ compatible with tree T has the following **Markov property**: if v and v' are children of u in T , then

$$\mathbb{P}\{X_{T_v} = x, X_{T_{v'}} = x' | X_u = y\} = \mathbb{P}\{X_{T_v} = x | X_u = y\} \mathbb{P}\{X_{T_{v'}} = x' | X_u = y\}.$$

In other words, the subtrees induced by the children are conditionally independent given the parent; this follows directly from the definition of the Markov tree measure in (10).

²One can easily construct a breadth-first ordering on a given tree by ordering the nodes arbitrarily within each level and listing the levels in ascending order: $\text{lev}_T(1), \text{lev}_T(2), \dots$

2.3 Statement of result

Theorem 2.1. Let \mathcal{S} be a finite set and let $(X_i)_{1 \leq i \leq n}$, $X_i \in \mathcal{S}$ be a Markov tree process, defined by a tree $T = (V, E)$ and transition kernels $p_0, \{p_{uv}(\cdot | \cdot)\}_{(u,v) \in E}$. Define the (u, v) -**contraction coefficient** θ_{uv} by

$$\theta_{uv} = \max_{y, y' \in \mathcal{S}} \|p_{uv}(\cdot | y) - p_{uv}(\cdot | y')\|_{\text{TV}}. \quad (11)$$

Suppose $\max_{(u,v) \in E} \theta_{uv} \leq \theta < 1$ for some θ and $\text{wid}(T) \leq L$. Then for the Markov tree process X we have

$$\bar{\eta}_{ij} \leq (1 - (1 - \theta)^L)^{\lfloor (j-i)/L \rfloor} \quad (12)$$

for $1 \leq i < j \leq n$.

To cast (12) in more usable form, we first note that for $L \in \mathbb{N}$ and $k \in \mathbb{N}$, if $k \geq L$ then

$$\left\lfloor \frac{k}{L} \right\rfloor \geq \frac{k}{2L-1} \quad (13)$$

(we omit the elementary number-theoretic proof). Using (13), we have

$$\bar{\eta}_{ij} \leq \tilde{\theta}^{j-i}, \quad \text{for } j \geq i + L \quad (14)$$

where

$$\tilde{\theta} = (1 - (1 - \theta)^L)^{1/(2L-1)}.$$

The bounds in (3) and (4) are for different metric spaces and therefore not readily comparable (the result in (3) has the additional convexity assumption; see [4] for a discussion). For the case where (14) holds, Samson's bound [8] yields

$$\|\Gamma\|_2 \lesssim \frac{1}{1 - \tilde{\theta}^{\frac{1}{2}}}, \quad (15)$$

and the approximation

$$\|\Delta\|_\infty \lesssim \sum_{k=0}^{\infty} \tilde{\theta}^k = \frac{1}{1 - \tilde{\theta}} \quad (16)$$

holds trivially via (5). In the (degenerate) case where the Markov tree is a chain, we have $L = 1$ and therefore $\tilde{\theta} = \theta$; thus we recover the Markov chain concentration results in [3, 6, 8] and the approximations in (15,16) become precise inequalities.

Remark 2.2. The bounds in (15) and (16) are approximate because (14) does not hold for all $j > i$ but only starting with $j \geq i + L$. The difference between $(1 - (1 - \theta)^L)^{\lfloor (j-i)/L \rfloor} = 1$ and $\tilde{\theta}^{j-i}$ for $i < j < i + L$ is at most $1 - \tilde{\theta}^{L-1}$ and affects only a fixed finite number ($L - 1$) of entries in each row of Γ and Δ . Since $\|\cdot\|_2$ and $\|\cdot\|_\infty$ are continuous functionals, we are justified in claiming the approximate bound, which may be quantified if an application calls for it. The statements in (15) and (16) are only meant to convey an order of magnitude.

2.4 Proof of main result

The proof of Theorem 2.1 is combination of elementary graph theory and tensor algebra. We start with a graph-theoretic lemma:

Lemma 2.3. Let $T = ([n], E)$ be a tree and fix $1 \leq i < j \leq n$. Suppose $(X_i)_{1 \leq i \leq n}$ is a Markov tree process whose law \mathbb{P} on \mathcal{S}^n is compatible with T (in the sense of §2.2.2). Define the set

$$T_i^j = T_i \cap \{j, j+1, \dots, n\},$$

consisting of those nodes in the subtree T_i whose breadth-first numbering does not precede j . Then, for $y \in \mathcal{S}^{i-1}$ and $w, w' \in \mathcal{S}$, we have

$$\eta_{ij}(y, w, w') = \begin{cases} 0 & T_i^j = \emptyset \\ \eta_{ijo}(y, w, w') & \text{otherwise,} \end{cases} \quad (17)$$

where j_0 is the minimum (with respect to $<$) element of T_i^j .

Remark 2.4. This lemma tells us that when computing η_{ij} it is sufficient to restrict our attention to the subtree induced by i .

Proof. The case $j \in T_i$ implies $j_0 = j$ and is trivial; thus we assume $j \notin T_i$. In this case, the subtrees T_i and T_j are disjoint. Putting $\bar{T}_i = T_i \setminus \{i\}$, we have by the Markov property,

$$\mathbb{P}\{X_{\bar{T}_i} = x_{\bar{T}_i}, X_{T_j} = x_{T_j} \mid X_1^i = yw\} = \mathbb{P}\{X_{\bar{T}_i} = x_{\bar{T}_i} \mid X_i = w\} \mathbb{P}\{X_{T_j} = x_{T_j} \mid X_1^{i-1} = y\}.$$

Then from (2) and (7), and by marginalizing out the X_{T_j} , we have

$$\begin{aligned} \eta_{ij}(y, w, w') &= \frac{1}{2} \sum_{x_j^n} |\mathbb{P}\{X_j^n = x_j^n \mid X_1^i = yw\} - \mathbb{P}\{X_j^n = x_j^n \mid X_1^i = yw'\}| \\ &= \frac{1}{2} \sum_{x_{T_i^j}} \left| \mathbb{P}\{X_{T_i^j} = x_{T_i^j} \mid X_i = w\} - \mathbb{P}\{X_{T_i^j} = x_{T_i^j} \mid X_i = w'\} \right|. \end{aligned}$$

If $T_i^j = \emptyset$ then obviously $\eta_{ij} = 0$; otherwise, $\eta_{ij} = \eta_{ijo}$, since j_0 is the “first” element of T_i^j . \square

Next we develop some basic results for tensor norms; recall that unless specified otherwise, the norm used in this paper is the total variation norm defined in (7). If \mathbf{A} is an $M \times N$ column-stochastic matrix: ($\mathbf{A}_{ij} \geq 0$ for $1 \leq i \leq M, 1 \leq j \leq N$ and $\sum_{i=1}^M \mathbf{A}_{ij} = 1$ for all $1 \leq j \leq N$) and $\mathbf{u} \in \mathbb{R}^N$ is *balanced* in the sense that $\sum_{j=1}^N \mathbf{u}_j = 0$, we have, by the Markov contraction lemma ([3], Lemma B.1),

$$\|\mathbf{A}\mathbf{u}\| \leq \|\mathbf{A}\| \|\mathbf{u}\|, \quad (18)$$

where

$$\|\mathbf{A}\| = \max_{1 \leq j, j' \leq N} \|\mathbf{A}_{*,j} - \mathbf{A}_{*,j'}\|, \quad (19)$$

and $\mathbf{A}_{*,j} \equiv \mathbf{A}[\cdot, j]$ denotes the j^{th} column of \mathbf{A} . An immediate consequence of (18) is that $\|\cdot\|$ satisfies

$$\|\mathbf{AB}\| \leq \|\mathbf{A}\| \|\mathbf{B}\| \quad (20)$$

for column-stochastic matrices $\mathbf{A} \in \mathbb{R}^{M \times N}$ and $\mathbf{B} \in \mathbb{R}^{N \times P}$.

Remark 2.5. Note that if \mathbf{A} is a column-stochastic matrix then $\|\mathbf{A}\| \leq 1$, and if additionally \mathbf{u} is balanced then $\mathbf{A}\mathbf{u}$ is also balanced.

If $\mathbf{u} \in \mathbb{R}^M$ and $\mathbf{v} \in \mathbb{R}^N$, define their tensor product $\mathbf{w} = \mathbf{v} \otimes \mathbf{u}$ by

$$\mathbf{w}_{(i,j)} = \mathbf{u}_i \mathbf{v}_j,$$

where the notation $(\mathbf{v} \otimes \mathbf{u})_{(i,j)}$ is used to distinguish the 2-tensor \mathbf{w} from an $M \times N$ matrix. The tensor \mathbf{w} is a vector in \mathbb{R}^{MN} indexed by pairs $(i,j) \in [M] \times [N]$; its norm is naturally defined to be

$$\|\mathbf{w}\| = \frac{1}{2} \sum_{(i,j) \in [M] \times [N]} |\mathbf{w}_{(i,j)}|. \quad (21)$$

The following “tensorizing” lemma will play a key role in deriving our bound (we suppress the boldfaced vector notation for readability):

Lemma 2.6. *Consider two finite sets \mathcal{X}, \mathcal{Y} , with probability measures p, p' on \mathcal{X} and q, q' on \mathcal{Y} . Then*

$$\|p \otimes q - p' \otimes q'\| \leq \|p - p'\| + \|q - q'\| - \|p - p'\| \|q - q'\|. \quad (22)$$

Remark 2.7. Note that $p \otimes q$ is a 2-tensor in $\mathbb{R}^{\mathcal{X} \times \mathcal{Y}}$ and a probability measure on $\mathcal{X} \times \mathcal{Y}$.

Proof. Fix q, q' and define the function

$$F(u, v) = \sum_{x \in \mathcal{X}} |u_x - v_x| + \|q - q'\| \left(2 - \sum_{x \in \mathcal{X}} |u_x - v_x| \right) - \sum_{x \in \mathcal{X}, y \in \mathcal{Y}} |u_x q_y - v_x q'_y|$$

over the convex polytope $U \subset \mathbb{R}^{\mathcal{X}} \times \mathbb{R}^{\mathcal{X}}$,

$$U = \left\{ (u, v) : u_x, v_x \geq 0, \sum u_x = \sum v_x = 1 \right\};$$

note that proving the claim is equivalent to showing that $F \geq 0$ on U .

For any $\sigma \in \{-1, +1\}^{\mathcal{X}}$, let

$$U_\sigma = \{(u, v) \in U : \operatorname{sgn}(u_x - v_x) = \sigma_x\};$$

note that U_σ is a convex polytope and that $U = \bigcup_{\sigma \in \{-1, +1\}^{\mathcal{X}}} U_\sigma$.³

Pick an arbitrary $\tau \in \{-1, +1\}^{\mathcal{X} \times \mathcal{Y}}$ and define

$$F_\sigma(u, v) = \sum_x \sigma_x (u_x - v_x) + \|q - q'\| \left(2 - \sum_x \sigma_x (u_x - v_x) \right) - \sum_{x,y} \tau_{xy} (u_x q_y - v_x q'_y)$$

over U_σ . Since $\sigma_x (u_x - v_x) = |u_x - v_x|$ and τ_{xy} can be chosen (for any given u, v, q, q') so that $\tau_{xy} (u_x q_y - v_x q'_y) = |u_x q_y - v_x q'_y|$, the claim that $F \geq 0$ on U will follow if we can show that $F_\sigma \geq 0$ on U_σ .

Observe that F_σ is affine in its arguments (u, v) and recall that an affine function achieves its extreme values on the extreme points of a convex domain. Thus to verify that $F_\sigma \geq 0$ on U_σ , we need only check the value of F_σ on the extreme points of U_σ . The extreme points of U_σ are pairs (u, v) such that, for some $x', x'' \in \mathcal{X}$, $u = \delta(x')$ and $v = \delta(x'')$, where $\delta(x_0) \in \mathbb{R}^{\mathcal{X}}$ is given by $[\delta(x_0)]_x = \mathbb{1}_{\{x=x_0\}}$. Let (\hat{u}, \hat{v}) be an extreme point of U_σ . The case $\hat{u} = \hat{v}$ is trivial, so assume $\hat{u} \neq \hat{v}$. In this case, $\sum_{x \in \mathcal{X}} \sigma_x (\hat{u}_x - \hat{v}_x) = 2$ and

$$\begin{aligned} \left| \sum_{x \in \mathcal{X}, y \in \mathcal{Y}} \tau_{xy} (\hat{u}_x q_y - \hat{v}_x q'_y) \right| &\leq \sum_{x \in \mathcal{X}, y \in \mathcal{Y}} |\hat{u}_x q_y - \hat{v}_x q'_y| \\ &\leq 2. \end{aligned}$$

This shows that $F_\sigma \geq 0$ on U_σ and completes the proof. \square

³We define $\operatorname{sgn}(z) = \mathbb{1}_{\{z \geq 0\}} - \mathbb{1}_{\{z < 0\}}$. Note that the constraint $\sum_{x \in \mathcal{X}} u_x = \sum_{x \in \mathcal{X}} v_x = 1$ forces $U_\sigma = \{(u, v) \in U : u_x = v_x\}$ when $\sigma_x = +1$ for all $x \in \mathcal{X}$ and $U_\sigma = \emptyset$ when $\sigma_x = -1$ for all $x \in \mathcal{X}$. Both of these cases are trivial.

To develop a convenient tensor notation, we will fix the index set $V = \{1, \dots, n\}$. For $I \subset V$, a tensor indexed by I is a vector $\mathbf{u} \in \mathbb{R}^{\mathcal{S}^I}$. A special case of such an I -tensor is the product $\mathbf{u} = \bigotimes_{i \in I} \mathbf{v}^{(i)}$, where $\mathbf{v}^{(i)} \in \mathbb{R}^{\mathcal{S}}$ and

$$\mathbf{u}[x_I] = \prod_{i \in I} \mathbf{v}^{(i)}[x_i]$$

for each $x_I \in \mathcal{S}^I$. To gain more familiarity with the notation, let us write the total variation norm of an I -tensor:

$$\|\mathbf{u}\| = \frac{1}{2} \sum_{x_I \in \mathcal{S}^I} |\mathbf{u}[x_I]|. \quad (23)$$

In order to extend Lemma 2.6 to product tensors, we will need to define the function $\alpha_k : \mathbb{R}^k \rightarrow \mathbb{R}$ and state some of its properties:

Lemma 2.8. Define $\alpha_k : \mathbb{R}^k \rightarrow \mathbb{R}$ recursively as $\alpha_1(x) = x$ and

$$\alpha_{k+1}(x_1, x_2, \dots, x_{k+1}) = x_{k+1} + (1 - x_{k+1})\alpha_k(x_1, x_2, \dots, x_k). \quad (24)$$

Then

- (a) α_k is symmetric in its k arguments, so it is well-defined as a mapping

$$\alpha : \{x_i : 1 \leq i \leq k\} \mapsto \mathbb{R}$$

from finite real sets to the reals

- (b) α_k takes $[0, 1]^k$ to $[0, 1]$ and is monotonically increasing in each argument on $[0, 1]^k$
- (c) If $B \subset C \subset [0, 1]$ are finite sets then $\alpha(B) \leq \alpha(C)$
- (d) $\alpha_k(x, x, \dots, x) = 1 - (1 - x)^k$
- (e) if B is finite and $1 \in B \subset [0, 1]$ then $\alpha(B) = 1$.
- (f) if $B \subset [0, 1]$ is a finite set then $\alpha(B) \leq \sum_{x \in B} x$.

Remark 2.9. In light of (a), we will use the notation $\alpha_k(x_1, x_2, \dots, x_k)$ and $\alpha(\{x_i : 1 \leq i \leq k\})$ interchangeably, as dictated by convenience.

Proof. Claims (a), (b), (e), (f) are straightforward to verify from the recursive definition of α and induction. Claim (c) follows from (b) since

$$\alpha_{k+1}(x_1, x_2, \dots, x_k, 0) = \alpha_k(x_1, x_2, \dots, x_k)$$

and (d) is easily derived from the binomial expansion of $(1 - x)^k$. \square

The function α_k is the natural generalization of $\alpha_2(x_1, x_2) = x_1 + x_2 - x_1 x_2$ to k variables, and it is what we need for the analogue of Lemma 2.6 for a product of k tensors:

Corollary 2.10. Let $\{\mathbf{u}^{(i)}\}_{i \in I}$ and $\{\mathbf{v}^{(i)}\}_{i \in I}$ be two sets of tensors and assume that each of $\mathbf{u}^{(i)}, \mathbf{v}^{(i)}$ is a probability measure on \mathcal{S} . Then we have

$$\left\| \bigotimes_{i \in I} \mathbf{u}^{(i)} - \bigotimes_{i \in I} \mathbf{v}^{(i)} \right\| \leq \alpha \left\{ \|\mathbf{u}^{(i)} - \mathbf{v}^{(i)}\| : i \in I \right\}. \quad (25)$$

Proof. Pick an $i_0 \in I$ and let $\mathbf{p} = \mathbf{u}^{(i_0)}$, $\mathbf{q} = \mathbf{v}^{(i_0)}$,

$$\mathbf{p}' = \bigotimes_{i_0 \neq i \in I} \mathbf{u}^{(i)}, \quad \mathbf{q}' = \bigotimes_{i_0 \neq i \in I} \mathbf{v}^{(i)}.$$

Apply Lemma 2.6 to $\|\mathbf{p} \otimes \mathbf{q} - \mathbf{p}' \otimes \mathbf{q}'\|$ and proceed by induction. \square

Our final generalization concerns linear operators over I -tensors. An I, J -matrix \mathbf{A} has dimensions $|\mathcal{S}^J| \times |\mathcal{S}^I|$ and takes an I -tensor \mathbf{u} to a J -tensor \mathbf{v} : for each $y_J \in \mathcal{S}^J$, we have

$$\mathbf{v}[y_J] = \sum_{x_I \in \mathcal{S}^I} \mathbf{A}[y_J, x_I] \mathbf{u}[x_I], \quad (26)$$

which we write as $\mathbf{A}\mathbf{u} = \mathbf{v}$. If \mathbf{A} is an I, J -matrix and \mathbf{B} is a J, K -matrix, the matrix product \mathbf{BA} is defined analogously to (26).

As a special case, an I, J -matrix might factorize as a tensor product of $|\mathcal{S}| \times |\mathcal{S}|$ matrices $\mathbf{A}^{(i,j)} \in \mathbb{R}^{\mathcal{S} \times \mathcal{S}}$. We will write such a factorization in terms of a bipartite graph⁴ $G = (I + J, E)$, where $E \subset I \times J$ and the factors $\mathbf{A}^{(i,j)}$ are indexed by $(i, j) \in E$:

$$\mathbf{A} = \bigotimes_{(i,j) \in E} \mathbf{A}^{(i,j)}, \quad (27)$$

where

$$\mathbf{A}[y_J, x_I] = \prod_{(i,j) \in E} \mathbf{A}_{y_J, x_i}^{(i,j)}$$

for all $x_I \in \mathcal{S}^I$ and $y_J \in \mathcal{S}^J$. The norm of an I, J -matrix is a natural generalization of the matrix norm defined in (19):

$$\|\mathbf{A}\| = \max_{x_I, x'_I \in \mathcal{S}^I} \|\mathbf{A}[\cdot, x_I] - \mathbf{A}[\cdot, x'_I]\| \quad (28)$$

where $\mathbf{u} = \mathbf{A}[\cdot, x_I]$ is the J -tensor given by

$$\mathbf{u}[y_J] = \mathbf{A}[y_J, x_I];$$

(28) is well-defined via the tensor norm in (23). Since I, J matrices act on I -tensors by ordinary matrix multiplication, $\|\mathbf{Au}\| \leq \|\mathbf{A}\| \|\mathbf{u}\|$ continues to hold when \mathbf{A} is a column-stochastic I, J -matrix and \mathbf{u} is a balanced I -tensor; if, additionally, \mathbf{B} is a column-stochastic J, K -matrix, $\|\mathbf{BA}\| \leq \|\mathbf{B}\| \|\mathbf{A}\|$ also holds. Likewise, since another way of writing (27) is

$$\mathbf{A}[\cdot, x_I] = \bigotimes_{(i,j) \in E} \mathbf{A}^{(i,j)}[\cdot, x_i],$$

Corollary 2.10 extends to tensor products of matrices:

Lemma 2.11. *Fix index sets I, J and a bipartite graph $(I + J, E)$. Let $\{\mathbf{A}^{(i,j)}\}_{(i,j) \in E}$ be a collection of column-stochastic $|\mathcal{S}| \times |\mathcal{S}|$ matrices, whose tensor product is the I, J matrix*

$$\mathbf{A} = \bigotimes_{(i,j) \in E} \mathbf{A}^{(i,j)}.$$

Then

$$\|\mathbf{A}\| \leq \alpha \left\{ \|\mathbf{A}^{(i,j)}\| : (i, j) \in E \right\}.$$

We are now in a position to state the main technical lemma, from which Theorem 2.1 will follow straightforwardly:

⁴Our notation for bipartite graphs is standard; it is equivalent to $G = (I \cup J, E)$ where I and J are always assumed to be disjoint.

Lemma 2.12. Let \mathcal{S} be a finite set and let $(X_i)_{1 \leq i \leq n}$, $X_i \in \mathcal{S}$ be a Markov tree process, defined by a tree $T = (V, E)$ and transition kernels p_0 , $\{p_{uv}(\cdot | \cdot)\}_{(u,v) \in E}$. Let the (u, v) -contraction coefficient θ_{uv} be as defined in (11).

Fix $1 \leq i < j \leq n$ and let $j_0 = j_0(i, j)$ be as defined in Lemma 2.3 (we are assuming its existence, for otherwise $\bar{\eta}_{ij} = 0$). Then we have

$$\bar{\eta}_{ij} \leq \prod_{d=\text{dep}_T(i)+1}^{\text{dep}_T(j_0)} \alpha \{ \theta_{uv} : v \in \text{lev}_T(d) \} \quad (29)$$

where $\text{dep}_T(\cdot)$ is defined in §2.2.1.

Proof. For $y \in \mathcal{S}^{i-1}$ and $w, w' \in \mathcal{S}$, we have

$$\eta_{ij}(y, w, w') = \frac{1}{2} \sum_{x_j^n} |\mathbb{P}\{X_j^n = x_j^n | X_1^i = yw\} - \mathbb{P}\{X_j^n = x_j^n | X_1^i = yw'\}| \quad (30)$$

$$= \frac{1}{2} \sum_{x_j^n} \left| \sum_{z_{i+1}^{j-1}} \left(\mathbb{P}\{X_{i+1}^n = z_{i+1}^{j-1} x_j^n | X_1^i = yw\} - \mathbb{P}\{X_{i+1}^n = [z_{i+1}^{j-1} x_j^n] | X_1^i = yw'\} \right) \right|. \quad (31)$$

Let T_i be the subtree induced by i and

$$Z = T_i \cap \{i+1, \dots, j_0-1\} \quad \text{and} \quad C = \{v \in T_i : (u, v) \in E, u < j_0, v \geq j_0\}. \quad (32)$$

Then by Lemma 2.3 and the Markov property, we get

$$\begin{aligned} \eta_{ij}(y, w, w') &= \\ &\frac{1}{2} \sum_{x[C]} \left| \sum_{x[Z]} \left(\mathbb{P}\{X[C \cup Z] = x[C \cup Z] | X_i = w\} - \mathbb{P}\{X[C \cup Z] = x[C \cup Z] | X_i = w'\} \right) \right| \end{aligned} \quad (33)$$

(the sum indexed by $\{j_0, \dots, n\} \setminus C$ marginalizes out).

Define $D = \{d_k : k = 0, \dots, |D|\}$ with $d_0 = \text{dep}_T(i)$, $d_{|D|} = \text{dep}_T(j_0)$ and $d_{k+1} = d_k + 1$ for $0 \leq k < |D|$. For $d \in D$, let $I_d = T_i \cap \text{lev}_T(d)$ and $G_d = (I_{d-1} + I_d, E_d)$ be the bipartite graph consisting of the nodes in I_{d-1} and I_d , and the edges in E joining them (note that $I_{d_0} = \{i\}$).

For $(u, v) \in E$, let $\mathbf{A}^{(u,v)}$ be the $|\mathcal{S}| \times |\mathcal{S}|$ matrix given by

$$\mathbf{A}_{x,x'}^{(u,v)} = p_{uv}(x | x')$$

and note that $\|\mathbf{A}^{(u,v)}\| = \theta_{uv}$. Then by the Markov property, for each $z[I_d] \in \mathcal{S}^{I_d}$ and $x[I_{d-1}] \in \mathcal{S}^{I_{d-1}}$, $d \in D \setminus \{d_0\}$, we have

$$\mathbb{P}\{X_{I_d} = z_{I_d} | X_{I_{d-1}} = x_{I_{d-1}}\} = \mathbf{A}^{(d)}[z_{I_d}, x_{I_{d-1}}],$$

where

$$\mathbf{A}^{(d)} = \bigotimes_{(u,v) \in E_d} \mathbf{A}^{(u,v)}.$$

Likewise, for $d \in D \setminus \{d_0\}$,

$$\begin{aligned} \mathbb{P}\{X_{I_d} = x_{I_d} \mid X_i = w\} &= \sum_{x'_{I_1}} \sum_{x''_{I_2}} \cdots \sum_{x^{(d-1)}_{I_{d-1}}} \\ &\quad \mathbb{P}\{X_{I_1} = x'_{I_1} \mid X_i = w\} \mathbb{P}\{X_{I_2} = x''_{I_2} \mid X_{I_1} = x'_{I_1}\} \cdots \\ &\quad \mathbb{P}\left\{X_{I_d} = x_{I_d} \mid X_{I_{d-1}} = x^{(d-1)}_{I_{d-1}}\right\} \\ &= (\mathbf{A}^{(d)} \mathbf{A}^{(d-1)} \cdots \mathbf{A}^{(d_1)})[x_{I_d}, w]. \end{aligned} \quad (34)$$

Define the (balanced) I_{d_1} -tensor

$$\mathbf{h} = \mathbf{A}^{(d_1)}[\cdot, w] - \mathbf{A}^{(d_1)}[\cdot, w'], \quad (35)$$

the $I_{|D|}$ -tensor

$$\mathbf{f} = \mathbf{A}^{(d_{|D|})} \mathbf{A}^{(d_{|D|-1})} \cdots \mathbf{A}^{(d_2)} \mathbf{h}, \quad (36)$$

and $C_0, C_1, Z_0 \subset \{1, \dots, n\}$:

$$C_0 = C \cap I_{\text{dep}_T(j_0)}, \quad C_1 = C \setminus C_0, \quad Z_0 = I_{\text{dep}_T(j_0)} \setminus C_0, \quad (37)$$

where C and Z are defined in (32). For readability we will write $\mathbb{P}(x_U \mid \cdot)$ instead of $\mathbb{P}\{X_U = x_U \mid \cdot\}$ below; no ambiguity should arise. Combining (33) and (34), we have

$$\eta_{ij}(y, w, w') = \frac{1}{2} \sum_{x_C} \left| \sum_{x_Z} (\mathbb{P}(x[C \cup Z] \mid X_i = w) - \mathbb{P}(x[C \cup Z] \mid X_i = w')) \right| \quad (38)$$

$$= \frac{1}{2} \sum_{x_{C_0}} \sum_{x_{C_1}} \left| \sum_{x_{Z_0}} \mathbb{P}(x[C_1] \mid x[Z_0]) \mathbf{f}[C_0 \cup Z_0] \right| \quad (39)$$

$$= \|\mathbf{Bf}\| \quad (40)$$

where \mathbf{B} is the $|\mathcal{S}^{C_0 \cup C_1}| \times |\mathcal{S}^{C_0 \cup Z_0}|$ column-stochastic matrix given by

$$\mathbf{B}[x_{C_0} \cup x_{C_1}, x'_{C_0} \cup x_{Z_0}] = \mathbb{1}_{\{x_{C_0} = x'_{C_0}\}} \mathbb{P}(x_{C_1} \mid x_{Z_0})$$

with the convention that $\mathbb{P}(x_{C_1} \mid x_{Z_0}) = 1$ if either of $\{Z_0, C_1\}$ is empty. The claim now follows by reading off the results previously obtained:

$$\begin{aligned} \|\mathbf{Bf}\| &\leq \|\mathbf{B}\| \|\mathbf{f}\| && \text{Eq. (7)} \\ &\leq \|\mathbf{f}\| && \text{Remark 2.5} \\ &\leq \|\mathbf{h}\| \prod_{k=2}^{|D|} \left\| \mathbf{A}^{(d_k)} \right\| && \text{Eqs. (20,36)} \\ &\leq \prod_{k=1}^{|D|} \alpha \{ \|\mathbf{A}^{(u,v)}\| : (u, v) \in E_{d_k} \} && \text{Lemma 2.11.} \end{aligned}$$

□

Proof of Theorem 2.1. We will borrow the definitions from the proof of Lemma 2.12. To upper-bound $\bar{\eta}_{ij}$ we first bound $\alpha \{ \|\mathbf{A}^{(u,v)}\| : (u, v) \in E_{d_k} \}$. Since

$$|E_{d_k}| \leq \text{wid}(T) \leq L$$

(because every node in I_{d_k} has exactly one parent in $I_{d_{k-1}}$) and

$$\|\mathbf{A}^{(u,v)}\| = \theta_{uv} \leq \theta < 1,$$

we appeal to Lemma 2.8 to obtain

$$\alpha \{\|\mathbf{A}^{(u,v)}\| : (u,v) \in E_{d_k}\} \leq 1 - (1 - \theta)^L. \quad (41)$$

Now we must lower-bound the quantity $h = \text{dep}_T(j_0) - \text{dep}_T(i)$. Since every level can have up to L nodes, we have

$$j_0 - i \leq hL$$

and so $h \geq \lfloor (j_0 - i)/L \rfloor \geq \lfloor (j - i)/L \rfloor$. \square

The calculations in Lemma 2.12 yield considerably more information than the simple bound in (12). For example, suppose the tree T has levels $\{I_d : d = 0, 1, \dots\}$ with the property that the levels are growing at most linearly:

$$|I_d| \leq cd$$

for some $c > 0$. Let $d_i = \text{dep}_T(i)$, $d_j = \text{dep}_T(j_0)$, and $h = d_j - d_i$. Then

$$\begin{aligned} j - i \leq j_0 - i &\leq c \sum_{d_i+1}^{d_j} k \\ &= \frac{c}{2}(d_j(d_j + 1) - d_i(d_i + 1)) \\ &< \frac{c}{2}((d_j + 1)^2 - d_i^2) \\ &< \frac{c}{2}(d_i + h + 1)^2 \end{aligned}$$

so

$$h > \sqrt{2(j - i)/c} - d_i - 1,$$

which yields the bound, via Lemma 2.8(f),

$$\bar{\eta}_{ij} \leq \prod_{k=1}^h \sum_{(u,v) \in E_k} \theta_{uv}. \quad (42)$$

Let $\theta_k = \max\{\theta_{uv} : (u,v) \in E_k\}$; then if $ck\theta_k \leq \beta$ holds for some $\beta \in \mathbb{R}$, this becomes

$$\begin{aligned} \bar{\eta}_{ij} &\leq \prod_{k=1}^h (ck\theta_k) \\ &< \prod_{k=1}^{\sqrt{2(j-i)/c}-d_i-1} (ck\theta_k) \\ &\leq \beta^{\sqrt{2(j-i)/c}-d_i-1}. \end{aligned} \quad (43)$$

This is a non-trivial bound for trees with linearly growing levels: recall that to bound $\|\Delta\|_\infty$ (5), we must bound the series

$$\sum_{j=i+1}^{\infty} \bar{\eta}_{ij}.$$

By the limit comparison test with the series $\sum_{j=1}^{\infty} 1/j^2$, we have that

$$\sum_{j=i+1}^{\infty} \beta^{\sqrt{2(j-i)/c} - d_i - 1}$$

converges for $\beta < 1$. Similar techniques may be applied when the level growth is bounded by other slowly increasing functions.

3 Discussion

We have presented a concentration of measure bound for Markov tree processes; to our knowledge, this is the first such result.⁵ In the simple case of the *contracting, bounded-width* Markov tree processes (i.e., those for which $\text{wid}(T) \leq L < \infty$ and $\sup_{u,v} \theta_{uv} \leq \theta < 1$), the bound takes on a particularly tractable form (12), and in the degenerate case $L = 1$ it reduces to the sharpest known bound for Markov chains. The techniques we develop extend well beyond the somewhat restrictive contracting-bounded-width case, as demonstrated in the calculation in (43).

The technical results in §2.4, particularly Lemma 2.6 and its generalizations, might be of independent interest. It is hoped that these techniques will be extended to obtain concentration bounds for larger classes of directed acyclic graphical models.

Acknowledgements

I thank John Lafferty and Kavita Ramanan for useful discussions and suggestions.

References

- [1] Amir Dembo, Peter Mörters, Scott Sheffield, “A large-deviation theorem for tree-indexed Markov chains.” <http://arxiv.org/abs/math.PR/0306045>, 2003
- [2] Leonid Kontorovich, “Measure Concentration of Hidden Markov Processes.” <http://arxiv.org/abs/math.PR/0608064>, 2006
- [3] Leonid Kontorovich and Kavita Ramanan, “Concentration Inequalities for Dependent Random Variables via the Martingale Method.” <http://arxiv.org/abs/math.PR/0609835>, 2006.
- [4] Leonid Kontorovich, “Metric and Mixing Sufficient Conditions for Concentration of Measure.” Paper in preparation, 2006.
- [5] Michel Ledoux, *The Concentration of Measure Phenomenon*, Mathematical Surveys and Monographs Vol. 89, American Mathematical Society, 2001.
- [6] Katalin Marton, “Bounding \bar{d} -distance by informational divergence: a method to prove measure concentration.” *Ann. Probab.*, Vol. 24, No. 2, 857–866, 1996.
- [7] Katalin Marton, “A measure concentration inequality for contracting Markov chains.” *Geom. Funct. Anal.*, Vol. 6, 556–571, 1997.
- [8] Paul-Marie Samson, “Concentration of measure inequalities for Markov chains and Φ -mixing processes.” *Ann. Probab.*, Vol. 28, No. 1, 416–461, 2000.

⁵In a 2003 paper, Dembo et al. [1] presented large deviation bounds for typed Markov trees, which is a more general class of processes than the Markov tree processes defined here. The techniques used and bounds obtained in [1] are of a rather different flavor than here; this is not surprising since measure concentration and large deviations, while pursuing similar goals, tend to use different methods and state results that are often not immediately comparable.